



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/715,225	11/17/2003	Arun Kwangil Iyengar	YOR920030488US1 (163-16)	5015
24336	7590	06/20/2006	EXAMINER	
KEUSEY, TUTUNJIAN & BITETTO, P.C. 20 CROSSWAYS PARK NORTH SUITE 210 WOODBURY, NY 11797			TSAI, SHENG JEN	
			ART UNIT	PAPER NUMBER
			2186	

DATE MAILED: 06/20/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No. 10/715,225	Applicant(s) IYENGAR ET AL.	
	Examiner Sheng-Jen Tsai	Art Unit 2186	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 17 November 2003.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-26 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-26 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. _____.
3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|---|---|
| 1) <input checked="" type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____ |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | 5) <input type="checkbox"/> Notice of Informal Patent Application (PTO-152) |
| 3) <input checked="" type="checkbox"/> Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)
Paper No(s)/Mail Date <u>11/17/2003</u> . | 6) <input type="checkbox"/> Other: _____ |

DETAILED ACTION

1. Claims 1-26 are presented for examination in this application (10,715,225) filed on November 17, 2003.

Acknowledge is made of information disclosure document filed November 17, 2003.

Claim Rejections - 35 USC § 102

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

2. Claims 1-6, 10-12, 16-24 and 26 are rejected under 35 U.S.C. 102(e) as being anticipated by lyengar et al. (U.S. 6,871,268).

As to claim 1, lyengar et al. disclose **in a system comprised of a plurality of storage elements** [figure 1 shows a system comprising a central cache (110), a remote server (104) and a plurality of processors (106-1~106-N) where the central cache and each of the processor has a copy of cache as storage element], **a method for maintaining objects in the storage elements** [Methods and Systems for Distributed Caching in Presence of Updates and in Accordance with Holding Times (abstract)] **comprising the steps of:**

maintaining information regarding which storage elements are storing particular objects [the central cache maintains local directories 110 which indicate the contents of local caches. A local directory maintains information about what objects may, but do not necessarily have to be, cached in the corresponding local cache. These local

Art Unit: 2186

directories 110 allow a central cache to update local caches (column 4, lines 30-35); the central cache 102 stores information from at least one remote server 104. The central cache communicates with a plurality of processors 106 which contain local caches 108. The central cache contains information about what is stored in local caches 108. When cached data changes, the central cache 102 is notified. The central cache is then responsible for updating local caches 108 (column 3, lines 59-65)] **in a consistency coordinator** [the central cache (figure 1, 110)] **which communicates with the storage elements** [the central cache 102 stores information from at least one remote server 104. The central cache communicates with a plurality of processors 106 which contain local caches 108. The central cache contains information about what is stored in local caches 108. When cached data changes, the central cache 102 is notified. The central cache is then responsible for updating local caches 108 (column 3, lines 59-65)];

responding to a request to update an object [in step 202 (figure 2), a request for an object is issued (column 5, lines 5-15)] **by using maintained information to determine which of the storage elements may store a copy of the object** [The central cache contains information about what is stored in local caches 108. When cached data changes, the central cache 102 is notified. The central cache is then responsible for updating local caches 108 (column 3, lines 59-65)];

instructing the storage elements, which the consistency coordinator suspects store a copy of the object, to invalidate their copy of the object [When cached data changes, the central cache 102 is notified. The central cache is then responsible

for updating local caches 108 (column 3, lines 59-65); it is to be understood that the term "update," as used herein, is meant not only to include changing the value of a data object in a cache but also may include invalidating the data object or performing some other operation on the object. The central cache may communicate remotely with processes running either on the same processing node or on different processing nodes. That way, several applications running on different processing nodes may communicate with the same cache (column 2, lines 15-24)]; **and performing an update of the object after each storage element that includes the copy of the object indicates that the storage element has invalidated the copy of the object or the storage element is determined to be unresponsive** [further, a local cache may exist on the same processing node as an application so that the application may obtain cached data locally. For cached objects, the central cache may keep a directory of which local caches are storing the object. Updates to cached objects may go through the central cache. In order to update a cache object, the central cache may communicate with the local caches to make sure that all copies are invalidated or updated (column 2, lines 25-32); in step 302 (figure 3), data changes, and the central cache is notified of the data changes. Notification may be by way of one or more of the processors 106 informing the central cache of any associated data changes. In step 304, the central cache coordinates cache updates. That is, the central cache updates all objects it has cached which have changed. In addition, the central cache consults its local directories 110 to see which local caches may contain changed

objects. Using local directories, the central cache 102 sends appropriate update messages to local caches (column 5, lines 40-50)].

As to claim 2, Iyengar et al. teach that **the step of maintaining information includes maintaining information regarding which storage elements are storing particular objects in the consistency coordinator** [The central cache contains information about what is stored in local caches 108. When cached data changes, the central cache 102 is notified. The central cache is then responsible for updating local caches 108 (column 3, lines 59-65); In addition, the central cache consults its local directories 110 to see which local caches may contain changed objects. Using local directories, the central cache 102 sends appropriate update messages to local caches (column 5, lines 40-50)].

As to claim 3, Iyengar et al. teach that **the consistency coordinator includes multiple nodes** [figure 1 shows a plurality of nodes of processors (106-1~106-N) where each of the processor has a copy of cache as storage element] **and each node of the consistency coordinator stores information for a different set of objects** [since local caches require extra space and may thus in some situations be of limited size, it is preferred to have one or more methods for determining which objects to store in a local cache. Such methods, referred to as cache replacement policies, are described below in accordance with the present invention (column 4, lines 25-30)].

As to claim 4, Iyengar et al. teach that **the storage elements include at least one cache** [figure 1 shows a plurality of nodes of processors (106-1~106-N) where each of the processor has a copy of cache as storage element].

As to claim 5, Iyengar et al. teach that **the storage elements are included in a distributed system** [figure 1 shows the configuration of a distributed system; in one aspect, a distributed caching technique of the invention comprises the use of a central cache and one or more local caches (column 2, lines 10-25)].

As to claim 6, Iyengar et al. teach **the method as recited in claim 1, further comprising the step of obtaining a lock on the object to be updated before performing the update** [with respect to the locking or holding time issue, in another aspect, the invention provides techniques for adaptively determining such time values (column 2, lines 39-51)].

As to claim 10, refer to "As to claim 1."

As to claim 11, refer to "As to claim 1."

As to claim 12, refer to "As to claim 3."

As to claim 16, refer to "As to claim 4."

As to claim 17, refer to "As to claim 1."

As to claim 18, refer to "As to claim 1."

As to claim 19, Iyengar et al. teach that **the system as recited in claim 18, further comprising a writer, which updates the object to be updated** [it is to be understood that the term "update," as used herein, is meant not only to include changing the value of a data object in a cache (column 2, lines 15-18). It is noted that changing the value of a data object inherently requires a write operation, hence a writer].

As to claim 20, Iyengar et al. teach that **the writer resides on a same node as a storage element** [for example, the central cache which would update the cache data objects as shown in figure 1].

As to claim 21, refer to "As to claim 1."

As to claim 22, refer to "As to claim 1."

As to claim 23, Iyengar et al. teach that **the system as recited in claim 18, further comprising at least one content provider** [for example, the central cache or the remote server as shown in figure 1; for instance, a cache may be implemented as a server in a network (e.g., a cache server or proxy caching server in a World Wide Web or Internet environment)(column 1, lines 18-20)].

As to claim 24, Iyengar et al. teach that **the content provider resides on a same node as a storage element** [for example, the central cache which would update the cache data objects as shown in figure 1].

As to claim 26, refer to "As to claim 4."

Claim Rejections - 35 USC § 102

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

3. Claims 1-5, 10-12, 16-18, 21-22 and 26 are rejected under 35 U.S.C. 102(b) as being anticipated by Hayes et al. (U.S. 5,497,480).

As to claim 1, Hayes et al. disclose **in a system comprised of a plurality of storage elements** [figure 1 shows a system comprising a plurality of nodes, each including a processor, a TLB (Translation Look-aside Buffer), a cache controller and cache memory (inside the cache controller block). A cache memory is one type of

storage element], **a method for maintaining objects in the storage elements**

[Broadcast Demap for Deallocating Memory Pages in a Multiprocessor System (title)]

comprising the steps of:

maintaining information regarding which storage elements are storing particular

objects [the TLB (Translation Look-aside Buffer) provides the information regarding

which “physical address” is associated with which storage element to support the

“virtual address to physical address translation” (column 1, lines 39-55)] **in a**

consistency coordinator [one of the plurality of processors (figure 1, 110, 120 or 130)

acts as the consistency coordinator] **which communicates with the storage**

elements [figure 1 shows that the plurality of processors (figure 1, 110, 120 or 130)

communicate with one another via the bus (figure 1, 100) by broadcasting a demap

request packet on the packet-switched bus (abstract)];

responding to a request to update an object by using maintained information to

determine which of the storage elements may store a copy of the object [page

table entries are removed from a plurality of TLBs in the multiprocessor computer

system by first broadcasting a demap request packet on the packet-switched bus in

response to one of the processors requesting that a page table entry be removed from

its associated TLB. The demap request packet includes a virtual address and context

information specifying this page table entry (abstract); when a page mapping of virtual-

to-physical addresses of a given process is swapped out or thrashed as the process

requires, the mapping has to be disposed of (column 1, lines 55-58); the method

comprises the steps of issuing a request packet by a first controller of the first TLB to

remove the page table entry from the first TLB; sending the request packet to the packet-switched bus to be broadcast to a second controller coupled to a second TLB with the request packet specifying a predetermined source, the first address mode and process identification; receiving the request packet by the second controller on the packet-switched bus; checking by the second controller to determine whether the second TLB contains the page table entry by comparing the first mode address and process identification; completing any pending operations for the second processor; removing the page table entry from the second TLB by the second controller if the page table entry is contained in the second TLB; issuing a reply packet by the second controller to indicate completion to the first controller; sending the reply packet to the packet-switched bus to be forwarded to the first controller with the reply packet identifying the source (column 2, lines 23-44)];

instructing the storage elements, which the consistency coordinator suspects store a copy of the object, to invalidate their copy of the object [controllers reply to the demap request packet by sending a first reply packet to the controller that sent the original demap request packet to indicate receipt of the demap request packet. If a controller removes the page table entry from its associated TLB, that controller sends a second demap reply packet to indicate that the page table entry has been removed from its associated TLB (abstract)]; and

performing an update of the object after each storage element that includes the copy of the object indicates that the storage element has invalidated the copy of the object or the storage element is determined to be unresponsive [the method

comprises the steps of issuing a request packet by a first controller of the first TLB to remove the page table entry from the first TLB; sending the request packet to the packet-switched bus to be broadcast to a second controller coupled to a second TLB with the request packet specifying a predetermined source, the first address mode and process identification; receiving the request packet by the second controller on the packet-switched bus; checking by the second controller to determine whether the second TLB contains the page table entry by comparing the first mode address and process identification; completing any pending operations for the second processor; removing the page table entry (i.e., invalidate the entry) from the second TLB by the second controller if the page table entry is contained in the second TLB; issuing a reply packet by the second controller to indicate completion to the first controller; sending the reply packet to the packet-switched bus to be forwarded to the first controller with the reply packet identifying the source (column 2, lines 23-44)].

As to claim 2, Hayes et al. teach that **the step of maintaining information includes maintaining information regarding which storage elements are storing particular objects in the consistency coordinator** [the TLB (Translation Look-aside Buffer) provides the information regarding which “physical address” is associated with which storage element to support the “virtual address to physical address translation” (column 1, lines 39-55); figure 1 shows that each processor/storage element has a copy of TLB; each processor/storage element may act as the consistency coordinator].

As to claim 3, Hayes et al. teach that **the consistency coordinator includes multiple nodes** [figure 1 shows a plurality of nodes of processors (110, 120 or 130)

where each of the processor has a copy of cache as storage element (i.e., cache memory)] **and each node of the consistency coordinator stores information for a different set of objects** [the TLB (Translation Look-aside Buffer) provides the information regarding which “physical address” is associated with which storage element to support the “virtual address to physical address translation” (column 1, lines 39-55); checking by the second controller to determine whether the second TLB contains the page table entry by comparing the first mode address and process identification; completing any pending operations for the second processor; removing the page table entry (i.e., invalidate the entry) from the second TLB by the second controller if the page table entry is contained in the second TLB (column 2, lines 23-44);].

As to claim 4, Hayes et al. teach that **the storage elements include at least one cache** [figure 1 shows a plurality of nodes of processors (110, 120 or 130) where each of the processor has a unit of cache controller (and the associated cache memory) as storage element].

As to claim 5, Hayes et al. teach that **the storage elements are included in a distributed system** [figure 1 shows the configuration of a distributed system comprising a plurality of processor/storage element].

As to claim 10, refer to “As to claim 1.”

As to claim 11, refer to “As to claim 1.”

As to claim 12, refer to “As to claim 3.”

As to claim 16, refer to “As to claim 4.”

As to claim 17, refer to "As to claim 1."

As to claim 18, refer to "As to claim 1."

As to claim 21, refer to "As to claim 1."

As to claim 22, refer to "As to claim 1."

As to claim 26, refer to "As to claim 4."

Claim Rejections - 35 USC § 103

4. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

5. Claims 7-9, 14-15 and 25 are rejected under 35 U.S.C. 103(a) as being unpatentable over Hayes et al. (U.S. 5,497,480), and in view of Chang et al. (US Patent Application Publication 2005/0128960).

As to claims 7-9, Hayes et al. do not mention that **sending heart beat messages to obtain availability information to and from a storage element**.

However, Chang et al disclose in their invention "Method for Determination of Remote Adapter and/or Node Liveness" a heart beat message protocol for the determination of node liveness in a distributed data processing system [abstract; figures 6-8; paragraph 0017].

Using hear beat messages allows early detections of any failure component and prompt recovery operations to maintain high availability of system [Chang et al., paragraph 0003].

Therefore it would have been obvious for persons of ordinary skills in the art at the time of the applicant's invention to recognize the benefits using hear beat messages to identify faulty components as soon as possible, as demonstrated by Chang et al., and to incorporate it into the existing apparatus and method disclosed by Hayes et al., to further improve the availability and reliability of the system.

As to claims 14-15, refer to "As to claims 7-9."

As to claim 25, refer to "As to claims 7-9."

6. *Related Prior Art of Record*

The following list of prior art is considered to be pertinent to applicant's invention, but not relied upon for claim analysis conducted above.

- Gannon et al., (US 5,265,232), "Coherency Control by Data Invalidation in Selected Processor Caches without Broadcasting to Processor Caches not Having the Data."
- Chang et al., (US 5,398,325), "Method and Apparatus for Improving Cache Consistency Using a Single Copy of a Cache tag Memory in Multi Processor Computer Systems."
- Butts, Jr. et al., (US 5,303,362), "Coupled Memory Multiprocessor Computer System Including Cache Coherency Management Protocols."
- Hayes et al., (US 6,073,212), "Reducing Bandwidth and Areas Needed for Non-Inclusive Memory Hierarchy by Using Dual Tags."
- McDonald et al., (US 6,012,127), "Multiprocessor Computing Apparatus with Optional Coherency Directory."

- Teramotop, (US 6,848,023), "Cache Directory Configuration Method and Information Processing Device."

Conclusion

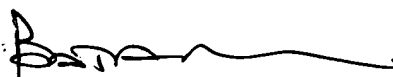
7. Claims 1-26 are rejected as explained above.
8. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Sheng-Jen Tsai whose telephone number is 571-272-4244. The examiner can normally be reached on 8:30 - 5:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Matthew Kim can be reached on 571-272-4182. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Sheng-Jen Tsai
Examiner
Art Unit 2186

February 14, 2006


PIERRE BATAILLE
PRIMARY EXAMINER
6/15/06